Transaction Cost, Transparency, and Innovation for the Internet.

By Shane Greenstein
Kellogg School of Management,
Northwestern University

FCC hearings, "Innovation, Investment, and the Open Internet"

January 13, 2009

What is it about the open Internet that facilitates innovation and investment? This essay will summarize several aspects of an answer to this question. This essay's answer involves two parts. First, this essay will focus attention on the importance of transaction cost for innovative behavior and outcomes. Second, the essay focuses on practices that shape transactions cost, namely, transparency and consistency, which, in turn, support wide participation in innovative activity.

This essay will not argue in favor of some specific legal or regulatory proposal. Rather, it will analyze these concepts and present a model for how they relate to one another. These concepts do have implications for the merits of specific regulatory and legal proposals, and some of these will be pointed out, but advancing such proposals is not the primary purpose.

The core thesis is as follows: The commercial Internet differs from its predecessors. The setting today necessitates some active managerial role for carriers. Concerns about the use of managerial discretion arise in the presence of market power. Concerns also arise as a result of the interdependence between carriers and other participants in the Internet value chain. Wide participation in innovative activity does not necessarily arise in the absence of a minimal level of transparency and consistency by certain vendors, such as platform providers and vendors with market power. That motivates a range of policy concerns.

The essay's thesis draws heavily from prior writing – most of it literally cuts and pastes from published articles. However, the essay will eschew the norms of legal scholarship and I will not cite scores of different sources. The references at the back of this essay list the three sources from which most of this essay's material is drawn. All of the key sources are in the footnotes of those three articles.

1. Transaction cost in collective invention

The value chain for Internet services appears to be perpetually in transition. That evolution raises a challenge for any regulatory framework: it makes it quite difficult to assess the general factors encouraging behavior that leads to innovative outcomes.

Transactions cost play an important role in the design and operation of that value chain. Transaction cost refers to two distinct areas of cost affiliated with two related activities. First, it refers to the cost of designing and setting up procedures to deliver a new service. Second, it refers to the cost of executing a set of proscribed processes and procedures of delivering a service to a user.

Almost by definition, transaction cost (in both senses) plays an important role because virtually every valuable activity on the Internet involves multiple participants – hardware vendors, software vendors, non-commercial participants, and users. Transactions cost arise whenever a participant tries to alter processes that multiple parties perform. Transaction costs arise whenever a participant tries to execute a set of processes that involve single or multiple parties.

Collective Invention

When focusing on innovative activity it is important to recognize two inventive eras for the Internet, one prior to its complete privatization by the NSF, roughly prior to the early 1990s. The other era is the commercial era with which we are familiar today. I will call the former pre-commercial and the latter commercial. The commercial era lives with some legacies of the pre-commercial era, and these legacies shape transactions cost.

In the pre-commercial era the Internet developed through a rather mundane process, accumulating capabilities over time from an enormous number of contributors. As such, it fits an archetype that scholars of innovation label as "collective Invention." Collective Invention is a "process in which improvements or experimental findings about a production process or a tool are regularly shared."[1]

[1] The quote is from Meyer (2003), who provides several modern examples that build the historical examples in Allen (1983).

For example, the creation, refinement, and improvement of e-mail prior to 1990 involved contributions from more than fifty different people over two decades, and that application was one new application among many.

By today's standards transaction cost in the pre-commercial era was high. Every new process required considerable effort to set up – namely, in comparison to today, where many aspects can be taken for granted. It involved notifying others. At the outset the community of inventors had to invent a process for doing that, in addition to inventing the process itself. Executing new processes involved high transaction cost because many parties had to learn to automate the coordination of multiple parties, such as backbone providers, network operators at a local level, and workstation users. By modern standards these too were high at the outset, and routines brought those costs down. By comparison today these costs are much lower, evident in the enormous scale of their use and execution today across a wide part of the globe.

The path we took does matter. It helps to explain why policy faces issues today.

Summarizing a long series of events, and simplifying greatly the extent of technical progress, the Internet initially accumulated capabilities over time in a government project hidden from mainstream view. Technical success generated interest and use, spread increasingly routinized processes among technically adept researchers, and gained economic value by growing capabilities in a community that did not (largely) recognize its economic value for non-researchers.

Once commercialized, the Internet began to accumulate more functions, as a range of firms began to use pieces of the Internet to enhance services. Over time the economic meaning of the Internet changed gradually, evolving from its technical roots to embody something greater than merely the technology. "The Internet" became a label for not only the Internet's core technical specifications, but for a wide range of economic activities – for an industry build to operate it and built upon its functions. "The Internet" stood for all the applications that accumulated around the Internet (such as the World Wide Web), used pieces of the Internet (such as many commercial and non-commercial applications), and developed new functions or delivered an enormous array of services to a wide range of users.

As a result, this essay employs a broad economic definition for the Internet, not solely a technical or engineering-oriented definition. And, consistent with that definition, it focuses on the

economic activities related to the Internet in this sense, namely, its contracts, it business conduct, and its firm structure.

As with the pre-commercial era, in the commercial Internet transaction cost played a role. Even the simplest of activities in this value chain, such as sending e-mail, involves many participants, and efficient delivery of services depends on advanced agreement about how their business activities will interrelate. However, transactions cost (in both senses) were much lower, and that supported a much different scale and set of activities. Today any discussion of transaction cost must account for additional factors, those that shape transparency.

Transparency

Transparent processes are those in which participants know what change is imminent. Participants in transparent processes inform others openly and vocally. In other words, participants' actions make it known—sometimes well in advance—when their changes will diminish or enhance the returns on others' innovative investments. In addition, participants can acquire information about others, and use that without restriction or limitation (e.g., they do not have to keep it secret).

In the pre-commercial era virtually all activity was transparent in principle (albeit, not necessarily in practice). That is, since most infrastructure design and application development took place under the broad sponsorship of either DARPA or NSF, there was an understanding of shared purpose, and expectation that researchers would let others know about their experiments. In addition, and perhaps more mundane in detail, there was a process for sharing information using the system based on Request For Comment (RFC), as sponsored by the Internet Engineering Task Force (IETF) from the late 1980s onward, and, more informally, by its predecessors in the academic community.

The IETF's processes still operate today, and today's Internet involves much more than just this group's activities. In the commercial era, the Internet inherited the norms of the commercial computing and telecommunications markets. In neither of those markets do suppliers all share a norm of transparency as a matter of principle, let alone practice. That said, in practice there tends to be a bit of transparency in many key activities for a variety of reasons.

To begin to illustrate this variety, let's first focus on one area in which transparency does play a role in the Internet. Transparency plays an important role in standards processes. Standards processes play an important role in shaping the transaction cost of setting up new processes and protocols.

Many participants in standards processes participate because they believe transparency has great importance in interdependent value chains. Other firms will not make long-term investments if they cannot understand at a fine level of detail how their software must interact with another firm's software or hardware.

As a result of transparent standards processes, here is one economic archetype for what happens after the issuance of a standard: Interested parties monitor the designs (because they can), and know that their near rivals do the same (because the data are available to anyone). Then all those parties plan to match each other along the dimension of the standard and differentiate along the dimensions in which each has competitive advantage (such as cost, features, brand, installed base, or distribution channel). Competition ensues once the standard is upgraded from its Beta to an endorsed and official standard.

In addition to transparency, several other factors play a role in shaping transaction cost. Next I focus on the role of platforms.

2. Platforms

The pre-commercial Internet contained carriers and applications and switches and hubs, as well as many other elements found today in the commercial Internet. Yet, the value chain has changed considerably. How so? One important change is associated with the rise of platform providers.

Well-designed standards and platforms hold one of the keys to the successful accumulation of functionality over time. Bresnahan and Greenstein (1997) state: "A computing platform is a reconfigurable base of compatible components on which users build applications. Platforms are most readily identified with their technical standards, i.e., engineering specifications for compatible hardware and software."

Platforms have become a central feature of the commercial Internet for facilitating such functionality. As noted, typical use of Internet-related services requires successful execution of a set of

technically interrelated activities coming from many independent firms. The failure or reduction in performance of any of these activities can lead to inferior outcomes for many users. Focusing solely on such technical action, however, misses a key dimension of how firms address the challenges.

To reduce the uncertainty about how such services interoperate, commercial firms take one of two approaches: either they negotiate arrangements in advance with all relevant participants, such as through a standard committee or consortia, or, if that fails, they do much of it themselves. In the parlance of business language, firms either negotiate standards with others so the task performs smoothly, or they offer a platform that accomplishes the task.

Platform strategies played an important role in computing before the commercialization of the Internet. For many reasons many firms organized their strategic approach for commercial opportunities on the Internet with similar approaches.

Ever since the emergence of the Internet, several leading businesses organized different platforms to alter the potential value chains for users and developers. There are proprietary and non-proprietary platforms, all of which interoperate to provide services and some of which compete at the same time.

The rise of platforms on the Internet is a source of both celebration and consternation. While platforms perform functions that firms and/or users value, their presence usually suggests that some firms/users are better off with them than without. At the same time, large or dominant platform leaders (usually) possess market power and some possess non-transparent processes. The presence of market power will raise questions about whether those firms use their discretion in ways that lead to more innovation.

The list of platforms today is long, but this review is necessarily brief. To illustrate the observation, I highlight three proprietary platform providers, Microsoft and Intel and Google. I will also make remarks about an alternative organizing form that does not involve proprietary platforms.

Perhaps the best known of the commercial platform providers is Microsoft, which develops and sells an operating system branded as Windows. It organizes the computing platform around the personal computer, as well as many Intel-based servers. To produce and deliver this product Microsoft engages with a multiplicity of actors, users (e.g., businesses and households), original equipment

manufacturers (OEMs, e.g., Dell, HP, and others), and application developers (e.g., software vendors). The operating system allows them to interact with one another for more efficient delivery of services.

Microsoft's platform strategy involves proprietary software. It stands at the middle of a large ecosystem, interacting with a range of firms, providing leadership that drives towards a wide range of applications. Specifically, Microsoft gives considerable and extensive technical information to applications developers about its APIs. At the same time, its source code always was and is a trade secret. Accordingly, Microsoft always holds back information, which, almost by definition, means it is not entirely transparent.

Microsoft is far from alone in such practices. That was the norm for the day when Microsoft founded, and, frankly, it is the norm at other companies today, such as Apple.

Microsoft today offers a mix of supporting functionality for the Internet, as it has for some time. For example, in the early 1990s it offered TCP/IP compatibility in Windows as means to enhance the features of its networking software. In the mid 1990s it offered a browser. It also has a search engine, and its live game platforms based on Xbox generates a significant amount of Internet gaming traffic. It also participates in numerous standardization efforts for applications that are common on the Internet, with a varying degree of resource commitments that would be difficult to summarize in any simple sentence.

Intel is another prominent platform provider. It too stands at the middle of a large ecosystem, interacting with a wide range of firms, providing leadership that drives towards the standard hardware design and specification used in most desk top computers, lap tops, and net books. It too tends to have limited transparency, keeping many trade-secrets to itself, almost by definition, and, like Microsoft, releasing loads of information about the technical specifications others need to interact to with Intel's products.

It behaviour differs from Microsoft's for a variety of reasons, but a simple one stands out; Intel interacts much more with hardware than software firms. The latter motivated Intel to invest in a wide range of activities, some of them far afield from microprocessor manufacturing. Some of these have drawn Intel into standards efforts and other activities that require it to be very forthcoming with information. Once again, the variety of commitments is rather difficult to summarize in any simple sentence.

For example, Intel helped design and sponsor USB and USB standards, including funding the testing for conformance. It also branched into sponsoring a Wi-Fi standard for laptops under the Centrino label, helping to design further upgrades to the underlying technical standard, which was designed by IEEE committee 802.11, and helping to fund conformance-testing organizations as well. More recently, it has invested heavily in designing and supporting other 802 wireless standards, known as Wi-Max. In addition, Intel has worked hard to develop a position as a microprocessor provider for standard designs of smart phone devices.

As a third example, let me discuss Google. Today many observers believe Google has the most effective platform on the Internet. Its search engine is the most popular in English by far, as well as in many other languages. That supports a very lucrative ad-placement business.

Like any platform leader, Google supports a large ecosystem of complementary application developers. Many other firms also expend considerable resources optimizing their web pages to appear high on Google's search results. Altogether, like any important platform, Google's actions have become central to the economic prosperity of others. Some observers believe this will only continue, as its popularity will allow Google to develop a range of products supporting its search business.

In many of Google's services its transparency largely mimics the norms of open source. For example, if Google can attract more developers by giving away information about a beta service, such as Android, then they do. The company also sponsors a number of programs to design data standards for Internet users that let them face lower switching cost between applications.

It is incorrect, however, to call the firm uniformly transparent. Google does not reveal how the details of page-rank work except at a broad level, or how its keyword auction ranks the quality of distinct bids, again, except at a broad level. They never will reveal the details, since both lie at the core of AdSense and Adwords, and revealing too much detail would not be in their economic interest. Now that Google's ambitions have started to take it into some content that compete with other firms, questions arise about the neutrality of its proprietary processes. Thus, I believe it is inevitable that Google will face tension about the transparency of these aspects of its platform.

As noted, many other prominent platforms are involved in providing service on to Internet users. These include Cisco (networking equipment), Research in Motion (Blackberry), Apple (iPhone, iPod), Yahoo! (search, news, mail), Oracle (enterprise databases), E-Bay (auctions), as well as many others. The above examples are only a few among many prominent commercial platforms shaping

development of the Internet. Each one of these platforms deserves a longer description, and the absence of that description is due to space-constraints, not their lack of importance.

Taking a "snapshot" of the economic structure today, the infrastructure supports a rather complex value chain involving the interoperability of many different commercial platforms. This is an enormous evolution. The present arrangement looks nothing like the Internet of the early 1990s, when it first commercialized. Once again, it suggests our notions about the Internet should evolve too, to consider approaches appropriate for this commercial environment.

More to the point, aphoristically, the Internet has been called a "network of networks" since it first began to diffuse out of its non-commercial origins. Yet, distilling the Internet to that aphorism is misleading about its structure today; it does not reflect how commercial behavior shaped the evolution of how the Internet gets used in the last decade and a half, and it does not reflect the factors that shape the evolution of transactions cost on the Internet. Leading firms and their business partners view the commercial Internet through the same lens they view activities in the rest of computing. For them, the commercial Internet is a "network of platforms."

An additional layer of decision making

Platforms added an additional layer of decision making to the provision of services. That came with a benefit, to be sure. Platform leaders do try to lower transaction cost for their own services. Many platform leaders also lower transaction cost for their business partners, so platforms can smooth transactions between participants with long term relationships.

But platforms also come with some strings attached. The Internet involves a mix of firms using distinct norms for transparency. For profit firms take action to help their own platforms, and not that of others. Growth tends to agglomerate the successful platforms, but they also stand in the way of complementary entry which holds the potential to oppose the commercial interests of the present platform leaders.

It is possible to characterize some of this behaviour. There are a variety of forms for governing platforms, but most share these four functions:

- Designing a standard bundle of technical implementations that others used in their applications;

- Operating processes to alter those standards and informed others about those alterations;
- Establishing targets and roadmaps to coordinate developer and user investments;
- Providing tools and alternative forms of assistance to others who wanted to build applications using their technical standards.

It is possible to do all four with transparent process, and it is possible to do all four with processes that are proprietary. These functions look quite different if they emerge from one where a unilateral decision maker manages them, as with a proprietary firm (e.g. Microsoft, Apple), or one where a non-profit organizes them, as with standards committees (e.g., IEEE committee 802.11).

Perhaps one historical example can illustrate. There was no profit-oriented organization providing platform leadership for the commercial Internet in mid 1995, while there was a profit-oriented set of leaders for the PC.

Looking more deeply, there were meaningful differences in the processes for decision making behind the four functions. Two commercial firms in the PC market, Microsoft and Intel, retained and guarded their right to make unilateral decisions about the pervasive standards embedded within the platform. Microsoft's processes were proprietary. Only the platform leaders had unrestricted access to information. At the same time, Microsoft and Intel invested heavily and deliberately in designing a standard bundle, informing others about its features, and providing tools to use it.

The Internet at the time, in contrast, employed a consensus process for determining the design of pervasive standards and protocols. The predominant processes employed documented standards and did not restrict access to these documents or their use by any participant in the Internet. The investment in tools and planning was somewhat haphazard, but information about its design was transparent – namely, documented in a publically accessible place through routine processes, often with writing help from the IETF, with results easily obtainable by anyone. That supported widespread participation from a technical community of insiders, and, as noted by many observers, a rather dynamic process for upgrading the infrastructure by many parties acting independently, but building upon one another's innovation.

Each of these structures has strengths and weaknesses. There are on-going debates among analysts about what those are. Predictions about the death of any specific format for platforms have

been greatly exaggerated, as all types seem to exist in some form. For the near future, I believe we can expect competition and coexistence between them, as well as on-going tension about the economic boundaries between them.

One of the sources of tension relates to transparency. Typically non-proprietary platforms are organized with policies promoting transparency, that is, access to information and lack of restriction about its use. Typically proprietary platforms are organized with policies with less transparency for all, or, explicit differences in transparency for business partners and non-partners, where the latter get better treatment than the former. Who is favored with more transparency can depend on the platform provider's self-interest.

The practices of proprietary platform provider differ from the practices of the pre-commercial Internet, and which dominated during the era of collective invention. These changes frame questions about whether the changes improve matters or not. Once again, one's view on that question depends on one's view about the specific strengths and weaknesses of the alternative forms for organizing developments. The essay will say more about that below.

Before doing that, however, the essay must focus on another factor also has changed the operations of the commercial Internet, the emergence of dominant access provision in regional areas.

3. Broadband dominance and transparency

Broadband access firms have an important role to play, and contribute to the value enjoyed by many users. However, not only are most access providers accustomed to operating with norms of less transparency than found in the pre-commercial Internet, but some of them have market power in local markets. Those factors can shape the transactions cost of others. This is a potential source of concern.

Today, most urban households face a duopoly of wire-line choice: (1) an offering from a local cable franchise, and (2) an offering from a local telephone company. In some locations, they also may face options for wireless providers, which potentially may convert the duopoly into a more competitive supply. In many suburban areas (less dense settings) households face that duopoly or only one wire-line provider. One wire-line provider of services arises in many rural settings or isolated small cities, where

households lack alternatives to dial-up internet service except through a satellite provider and/or other wireless ISP.

As a cause for both celebration and concern, broadband firms inhabit a position of monopoly or duopoly in a key part of the value chain. Why celebrate this? Broadband's position reflects the ascendency of a superior product and service in replacing dial-up, which is an unambiguous economic improvement over the near past. After all, a decade ago fewer than 5% of US households had access to broadband. Why is it a source of consternation? It raises concerns about the presence of market power and its distortions on the incentives to innovate.

There is a traditional argument about the potential distortions market power puts on innovation. At a broad level, while society benefits from giving incentives to firms to create superior products and services, rewarding firms with monopoly power leads to high prices for their services while their provide those services. While this may be a source of concern, this essay will largely focus its attention elsewhere, on the factors that shape innovative activity.

Firms with market power may face weaker incentives to innovate than firms in any more competitive market structure. Fear of cannibalization and excessive institutional inertia around existing technological paradigms are the typical concerns. Both lead a firm insulated from competitive pressure to introduce new innovation more slowly (or not at all) than would occur in a competitive setting.

While those economic incentives would be present with or without transparency, and are potentially a source of concern, today I will focus primarily on how these interact with the factors shaping transaction costs. Once again, the theme will be a mix of celebration and consternation.

Specifically, the ascendency and diffusion of broadband enabled a range of applications to blossom, and in many instances the efforts and investments of broadband carriers played a positive role in that blossoming. Generally speaking, four types of rather different uses share the same capacity: (1) browsing and e-mail, which tend to employ low bandwidth and tolerate delay; (2) video downloading, which can employ high bandwidth and can tolerate some delay; (3) voice-over IP and video-talk, which tend to employ high bandwidth and whose quality declines with delay; and (4) peer-to-peer applications, which tend to use high bandwidth for sustained periods of time, and can tolerate delay, but, in some applications (*e.g.*, Bit-Torrent) can impose delay on others.

While that diversity of applications wrings additional productivity out of the same capital supporting the network, it comes with a potential drawback: the use of one application sometimes affects the productivity of another. In part this is due to capacity constraints at bottleneck positions in the network, or there are few backbone pathways to support browsing in isolated positions. Contributing to these constraints are geographically localized negative externalities – *e.g.*, many modern peer-to-peer applications employ all available bandwidth, diminishing the quality of other applications in the same cable network that cannot tolerate delay. In any case, improving efficiency requires some management of competing interests and users.

At present the actions of access providers is governed by rather light norms of transparency, such as the legal limitations inherent in service contracts with users. In the era of competitive provision of dial-up access this was of little concern. Most providers inherited the legacy norms and practices of the pre-commercial Internet, and users had many options to switch carriers. It is a bit of an exaggeration, but not far from the truth, that competitive forces prevented the emergence of bottlenecks and/or their exploitation, and there was little reason for concern.

The lack of transparency norms today is somewhat remarkable, especially in light of the complexity of some applications. In principle a carrier today can operate with almost no obligation to make their practices known to other participants.

Of course, practice differs for numerous reasons. Communication between firms is quite common, and so is access to the same set of engineering knowledge about how the network operates. In the recent past it was quite difficult for any firm to do anything novel (from an engineering level) without others knowing what was being done, knowing about it as it was being done, or finding out about it eventually. In addition, many of the largest broadband carriers in the US agreed to abide by the FCC's "four principles" as conditions for accepting mergers, and despite the ambiguity of those principles, these seemed to preclude many actions.

Not surprisingly, therefore, events that illustrate the broader issues occurred only recently. Much of the debate also concerns questions about how much to generalize from a small set of examples.

For example, transparency played a role in Comcast's unilateral declaration to throttle P2P applications on its lines with resets. There were many facets to this event, I will focus on aspects that

illustrate general points about the transparency of network management practices, as well as the transparency of other facets of firm conduct.

As it played out, one striking feature about this event was the willingness of all parties to act without telling anyone in advance what was happening. P2P users acted as if they could run any application on any time of day, irrespective of its consequences for others, even when it degraded the quality of service for neighbors during peak-load time periods. Comcast acted as if it had full discretion to manage its data over its facilities without informing any other suppliers of applications on the Internet, even its own customers, and, of course, the other users with whom Comcast's customers were communicating and sharing files.

This feature is even more striking because the carrier had a well understood goal, bringing some efficiency to the use of the assets shared by all users of Comcast's network. The economics were straightforward. Management could internalize the externality one user imposes on others—managing traffic for many users' general benefit. That is, P2P applications, like Bit-Torrent, can impose large negative externalities on other users, particularly in cable architectures during peak-load time periods.

Hence, on one level, Comcast's goals do not depart from widely accepted principles in a market economy – namely, a firm should try to manage their assets to reduce cost and enhance efficient use for a customer base. In addition, a firm should try to manage the quality of the experience for its own customer base, and make choices to trade-off one its own customer's quality with another.

Why not give Comcast unfettered discretion to manage the situation then? There is at least one additional transaction cost to consider, that between Comcast and other providers of applications, namely, other than Bit-torrent. That includes innovative entrepreneurs with plans to develop further applications, and it might include those who are not in the market at present, but might be in the near term. In addition, it might include other users, those who are not direct Comcast customers, but do communicate with Comcast customers, and anticipate certain operational practices.

Unfettered discretion for Comcast could raise the transaction cost to many other application developers, particularly if Comcast retains the right to remain non-transparent in its management policies. If Comcasts' policies about network management further remain shrouded, then a future entrant cannot develop applications without knowing what to expect, when it will change, and how it might be altered in the future. That is a high transactions cost for setting up innovative applications.

In that sense, Comcast's behavior had many less appealing aspects, such as its lack of transparency, as well as its virtually one-sided negotiating stance with all other application providers, and lack of clear statements about its own actions in advance or even while they were taking place.

As noted at the outset, it is not my goal to fully analyze the Comcast/Bit-Torrent events. I raised this example with a narrow goal in mind, to illustrate the role of transparency in broadband access and innovation. More to the point, this example raises at least two alarming general possibilities:

First, what if each of the major carriers in the United States used their discretion to pursue quite distinct approaches to managing their broadband operations, and did not make those practices transparent to other firms or other users? That would raise transactions cost for many other providers of new innovative applications, as each new change worked its way through a maze of a variety of technical issues, each different in different carrier's network, each subject to change without notice.

Second, what if carriers take other actions that shape the innovative returns of application providers, such as make investments in key technologies (or not), or adopt/alter policies for supporting their own services at the expense of others? That too would raise transactions cost for many other providers of new innovative applications.

In conclusion, it is not obviously healthy for the innovative activity of others to give Comcast's management (or any other firm's management) unchecked discretion to make take-it-or-leave-it offers to providers of any application its management believes harms users or its own interests.

4. Participation and Consistency differ from transparency

At several points in this discussion I have referred to others norms in practice that shape transaction cost. I now make these explicit by discussing norms for participation and consistency.

Participatory processes are those in which sponsoring organizations invite comment, discussion, and input from others affected by their actions. Organizations seeking wide participation solicit input through public forums, e-mail lists, blogs, community sites, and a range of other activities.

Standards organizations can illustrate. These organizations, such as the IEEE committees that determine the 802.11 standards, or the IETF, vary considerably in their policies for encouraging or

discouraging participation. For example, some organizations require fees, some require participants to meet certain technical qualifications, and others will allow any observer to attend, though not vote.

Wide participation is found in many standards processes, and quite frequently in open source projects, particularly those without sponsorship. Wikipedia, though non-profit, is perhaps the best-known example of an online project that encourages wide participation. The Firefox browser community, also non-profit, has quite diverse participation from numerous corners as well. So, too, does Linux, Apache, the IETF, and the World Wide Web Consortia. In the latter cases, most participants are quite technically skilled, so minimal skill levels can limit participation but appears not to so greatly.

Variety of participation norms is understandable. For example, wide participation is probably the least common attribute among standards consortia sponsored by commercial private firms. Most managers prefer to retain decision-making authority, guarding investment decisions in the name of stockholders. There is concern that giving up such discretion risks having participants take investment in directions that do not serve firm interests.

Accommodating wide participation in innovative activity is the central concern of this essay. That goal raises numerous issues related to transaction cost, and, as it turns out, no single answer suffices. In part, this is due to the drawbacks of wide participation. It may lead to slower decision making and more onerous managerial challenges coming to consensus. There are strengths and weaknesses to varying approaches, and no single approach dominates.

Consider, for example, the design of Web standards maintained by the World Wide Web Consortium (W3C). It has a transparent process, and it employs a decision making model with some restrictions on participation. The W3C requires firms to pay for their membership. Nonetheless, the consortia involves (and invites) views from a wide variety of users of the Web. At the same time, Tim Berners-Lee and his staff retain some authority to make decisions unilaterally after consultations with the membership, partly to avoid slowness and potential deadlocks.

Another open source community, Webkit, illustrates a different type of arrangement. After unilaterally altering the code base for several years and giving it back to the open source community, Apple then promulgated a licensing scheme for Webkit that enabled wide participation, using a variant of a widely used General Purpose License (GPL). The GPL seemed to persuade other firms and developers that Apple would not solely determine the direction of Webkit's applications, but, rather, a community of participants with potentially wide views would have rights to take the code in a variety of

direction (and split with Apple, if they chose to). Webkit now serves as the code base for the rendering engines in the browsers for many important non-Microsoft cell phones – the Safari browser by Apple, the Chrome browsers by Google, and the S60 Nokia browser, for example.

Another well known example of these open source platforms is Linux, and this illustrates another approach. The changes to this open source project also hint at how commercialization and open source have both changed. Linux began as a volunteer project by Linus Torvalds, but today has firm support for a consortium operated by Torvalds. He and others draw a salary. This consortium supports a range of businesses operated by many firms, including IBM, Red Hat, and others. It receives contributions from a wide set of participants too, some of them employed at these sponsoring firms and some from outside that group.

Many open source advocates argue that transparency goes hand in hand with wide and large number of participants in decision making for standards and protocols. That is not exactly precise enough for a policy discussion, however.

Certainly there is a grain of truth to what open source advocates say. An interdependent value chain requires a minimal level of transparency about the basic engineering of the Internet and its routine execution. Such minimal levels of transparency supports lowering the transactions cost of setting up innovative new services, which enables wide and large involvement in innovative activities, which established firms and entrepreneurs choose to undertake.

However, it is naïve to argue that transparency and wide participation comes for free. Transparency about the regular management of resources can be costly to those who are required to be transparent – management must notify others, communicate with questions, and make effort to make sure the right participants are involved. Similarly, managing a wide set of participants is also costly.

In short, there is a cost and a benefit. That supports an interesting set of questions about trade-offs. Transparency from one key firm can lower the cost of others who services depend on interacting with that key firm. Yet, that key firm may have to incur a cost to help all those others

Does a firm have proper incentives to incur such cost? In other words, if management were given discretion, would they make decisions about transparency that met societal aims?

I conjecture not, particularly in the presence of market power. I would conjecture that those with market power face incentives to not incur the cost of transparency that fully internalizes the gains

that others reap from such transparency. In part incentives are low because providing information yields gains for others than a firm simply does not internalize. In part incentives are low in the presence of market power because the fear of cannibalization will induce a firm to resist providing information to potential competitors, particularly other platform providers who may not compete today but might in the near term. In part incentives are too low because the size and scope on the Internet is so large, so the gains are wide and diffuse, well outside the range of feasible commercial contracts by a single firm.

Overall, therefore, rather straightforward economic reasoning suggests that broadband firms with market power face weak economic incentives to notify others about changes to their management practices. In addition, I conjecture that the cost of doing such activity is far lower than the benefits others gain from knowing about those practices, so the incentives are inappropriate. I also conjecture that the sum of the gains from lowering of transaction cost among many application and software providers far exceeds the cost of making the notification.

I would go further, that this is also so for proprietary platform providers for many of the same reasons just stated. However, one important qualification needs to be stressed.

Platform leaders often deliberately set up eco-systems of complementary relationships, designating some firms as direct partners. In that instance, there are strong economic incentives to support direct partners with more transparency than others – that is, to provide partners with more information about the operation of the platform, the recent directions of change, the long term plans for change, and so on. The incentives are strong because the platform provider directly sees the benefit in their own economic prosperity if they support their business partners so directly.

With proprietary platforms, therefore, there may be tradeoffs between the costs and benefits of transparency for a limited number of firms chosen by a platform firm, and an alternative setting where the costs and benefits of being transparent for a larger group.

Consistency

There is a movement among some platform providers today to adopt consistent policies, Microsoft most prominently. Not all have done this, but this topic is worthwhile to discuss because of the confusion it raises. This is not the same as transparency or wide participation.

Consistent policies from a firm are those that change slowly at most, allowing for the planning of others. They are changed without caprice, without an ad hoc approach, and without seemingly arbitrary timing, in other words, without actions that necessarily diminish the returns on others' innovative, long-term investments.

Concerns about consistency arise for many of the same reasons addressed in the discussion about transparency. It shapes the profitability and investments of others.

Partners and non-partners to platform providers worry about whether they receive the same treatment as others over time, and about whether platform providers will act in opportunistic ways – taking advantage of urgent moments to renegotiate a key price or contractual provision, for example. This can matter in an interdependent value chain for straightforward reasons: Application developers make investments whose value depends on how they operate with other complements. As they make these investments, want to predict the range of value they will generate value in the near future, and they also want to predict the range of factors that shape their ability to capture some of that value.

Consistency has an importance role in interdependent value chains, and especially networks. In such settings every application firm must work with every carrier.

Application firms will hesitate to make long-term investments if they fear not making a return on that investment due to changes by others, which are out of their control. Entrepreneurs will hesitate to take action if they fear conditions will change arbitrarily on them later, or systemically to their disadvantage. Firms will hesitate to undertake costly economic experiments if they cannot assure themselves that other firms won't interfere with the conditions that support learning from their market experience.

To be sure, many VCs and small start-ups care about such matters. Many fear expending scarce resources and managerial time on negotiation and frequent renegotiation, simply because it is costly. Many fear economic expropriation through inconsistent policies from commercial platform providers, namely, opportunistic renegotiation of contracting terms in the face of urgency, and/or unexpected events that alter the economic environment.

Microsoft recently and very publicly declared that it had adopted a set of principles that bound the firm to remain consistent in its actions over time. Not long after that first announcement Microsoft

announced another set of principles for remaining consistent in its interoperability designs, and these reinforced the earlier points.

Those two actions directly addressed one of the issues that perennially arose in the 1990s—accusations that Microsoft's employees altered APIs, contracting terms, or other firm technologies in self-interested ways that discriminated between business partners. This was thought to be a policy that application developers would find encouraging. It also relieved concerns about the potential waste of time, cost and effort out of negotiations.

In adopting a commitment to consistency, Microsoft did not give up its rights to retain secrets (e.g., remain less than transparent) nor to give up its right to retain managerial unilateral discretion (e.g., exclude participation from outsiders). Instead, Microsoft committed to not arbitrarily alter or apply what was decided unilaterally by management, inviting business partners (i.e., especially developers and OEMs) to inquire whether they receive treatment similar to another partner of Microsoft's (i.e., another developer's competitor).

Such a commitment represents a useful step towards allowing others in a complex value chain to plan, which can lower transaction cost for setting up new services. However, only a lawyer with a taste for nuance could find such a policy fully comforting if such a policy stood alone, unaccompanied by policies for transparency and participation. Frankly, as a practical matter, it is hard to discover inconsistency without full transparency.

In my experience with VCs and entrepreneurs I rarely see entrepreneurs or venture capitalists willingly roll the start-up dice for several years on a large investment based solely on such a promise about consistency. It is not equivalent to policies for supporting wide participation or transparency.

More to the point, contracting under legal norms should happen, but alone it does not seem to be sufficient in many circumstances. In my experience VCs and entrepreneurs prefer long experience with top managers at the platform firms to merely contracting alone. Experience supports deep understanding about senior management's trustworthiness, priorities, and willingness to settle differences in unexpected circumstances. If they have to solely rely on legal limits they are quite wary about finding themselves in a situation where they must enforce a legal provision in court over many years.

This is a good place to make a short digression. It is possible to make similar observations about the (in)consistency in government regulatory policy. Consistency can lower transactions cost. However, I think this is only a step, and for many of the same reasons above. It is not as essential as participation and transparency in regulatory decision making. Moreover, in regulatory processes lack of participation and transparency is a good method for hiding lack of consistency. End of digression.

These remarks about consistency relate to the last topic – contractual incompleteness over new contingencies. Most participants in the commercial Internet anticipate the market environment will change in unexpected ways, and the incompleteness of their present contacts exposes them to many potential challenges in the near term. Transparency and inconsistency play a role in how firms react to those challenges.

5. Contractual incompleteness and transactions cost

Contractual incompleteness has been a central feature of the Internet value chain since it commercialized. Incompleteness refers to the absence of contracts governing regular transactions or, if such contracts exist, to contracts that lack fully specified terms for all contingencies.

When the Internet first commercialized many facets of its value chain lacked contracts. Over time many firms designed new processes and activities, altered their conduct, and developed contractual relationships with other business partners, which filled in some of the incompleteness.

It is sometimes said that the maturation of the Internet value chain after its commercialization has diminished this incompleteness somewhat, but there are a range of opinions about how much. My own opinion is that the Internet is not mature yet, but is closer to mature than it was when it commercialized in 1995.

A mature value chain is one in which every aspect of its value chain has a lengthy history, and all alternatives arrangements are known or reasonably explored. In such a setting, the likelihood of dramatic change is low, and so too is the likelihood of contractual incompleteness.

I do not think the Internet has reached such a mature state yet, nor will it for some time. Too many new technologies in integrated circuits, switching, small devices, and cloud computing (to name just a few) continue to open new commercial opportunities. Too many established firms and VCs want

to invest in understanding and taking advantage of those opportunities. So, while the lack of maturity might not continue indefinitely, I think contractual incompleteness will continue to play an important role in the near term.

This observation is relevant to innovation policy. It provides both justification for government intervention, as well as a limitation to it.

Contractual incompleteness arises for many reasons. The Internet involves an extraordinarily large number of parties, which renders multi-lateral negotiations impractical in many settings, even in some standards committees (which can be a very good place to conduct such negotiations). There are so many players, in part, because the value chain supports an extraordinarily multi-purpose network, as earlier noted. Said simply, today many parties take action and their actions influence one another. There is just no practical way to get all these participants—or even their representatives—in the same room at the same time to work out a deal by horse-trading one set of economic concerns for another.

Incompleteness also arises where all parties may recognize the potential for technical change to generate new applications that alter circumstances, requiring renegotiation of prior contracts whose terms are no longer relevant. Yet, many pairs of parties in this setting may fail to come to agreement for numerous reasons. Even if the recognition exists, the parties may fail to negotiate a solution due to lack of the type of trust and mutual assumptions that usually supports renegotiating commercial transactions in the face of such contractual incompleteness.

Most interesting – and as already noticed – contractual incompleteness may arise because the relevant party may not even exist yet (if they will be entrepreneurial start-ups) and, thus, lack representation in even a basic form, such as trade-group or related commercial organization. That is an especially salient issue in industries, such as the Internet, which lack maturity. That is, decisions today shape entrepreneurship tomorrow, but few speak up for those future interests. Therefore, there is policy interest in protecting the conditions that support later entry of entrepreneurs, even though few today are there to keep established firms from raising transaction cost on later participants in the value chain.

Legal or regulatory ambiguities for innovative activities also can play a role. While contractual obligations govern some of the routine activities, it may be more difficult to erect similar obligations for new activities. For example, contracts govern the handoff of data from one backbone carrier to another,

or from one Web application to an edge-caching site, such as Akamai's, or to a content-delivery network, such as Amazon's.

In contrast, a looser contractual foundation governs another set of interrelated activities. For example, YouTube was founded in an era when there were multiple plausible definitions for a precise, legal, and safe-harbour for copyrighted material for user-supplied video. These definitions still remain ambiguous, though court cases continue to refine them into a less ambiguous domain.

Stated simply, the very thing that makes the Internet economically successful—the accumulation of innovation that supports a wide set of applications for many participants, including entrepreneurs—gives rise to conditions that make it harder to negotiate around the uncertainty. More to the point, while the value chain probably will look similar next month, only a naïve fool expects little change over several years, namely, the time periods over which many innovative investments reap their returns.

The failure of negotiations can provide an economic justification for a potential role for government regulators in specific circumstances: to settle disputes when many participants have a stake in the solution but private parties fail to account for these externalities; to mandate arbitration in cases when bilateral bargaining fails to reach resolution; or, related, to define "default" terms of commercial relationships that many partake in, when the default remains undefined; to mandate terms of standards employed by participants in the value chain when they otherwise cannot or do not come to such standards on their own; and, to reduce transaction cost for future entrepreneurs when established firms would prefer not to do so.

Note, however, this argument implies a limitation on that role. It covers only those activities that firms could not already settle themselves through contracting, those with large externalities, or those which necessarily involve unanticipated circumstances. To be sure, however, that is not necessarily a substantial limitation if it involves participants who are not even in a market yet, such as entrepreneurs.

Stated simply, negotiations happen, and so does renegotiation. It is endemic to the Internet value chain. Consequently, transaction cost shape negotiation, and are shaped by them. This is the last topic to discuss in this essay.

Negotiation and refusal to deal

In a network with a high degree of technical interrelatedness, there are general gains to all parties from bringing routines into business processes and activities, which lowers transaction cost, much like there are gains to adopting standards and platforms to coordinate activities. However, adopting such routines may require negotiation between multiple parties.

Such negotiation offers no guarantee of success. Many outcomes are possible. Occasionally both parties want an agreement, but just as often one party will desire it more than the other. Alternatively, one party may have an ability to generate a better deal than the other, and, thus, perceives moments of negotiation as an opportunity to generate a strategic advance or gain additional revenue. As a general rule, the structure of bargaining sometimes can work out to a Goldilocks equilibrium that is just right—not too hot and not too cold—but more often it does not. One firm gets too powerful or another prominent bargainer loses its way.

Let me define one term-of-art for this observation. In the extreme, bargaining becomes one-sided, with one party asking for something while the other refuses to provide it or only agrees to it at a high cost. The simplest manifestation of this extreme situation arises when the more powerful party declares a “take-it-or-leave-it” offer, leaving other parties no choice but one that favors the powerful party, or “refuses to deal”, leaving other parties with no choice at all, if the more powerful party perceives that no deal is in their interest.

The absence of one-sided bargaining and the absence of refusal-to-deal is a sign of well-functioning bargaining environment, while the presence of one-sided bargaining is a sign of potential illness, which might have adverse consequences that might spread. The key question is whether the less powerful parties have access to reasonable alternatives. This will take some explaining – in two steps, first on the role of negotiation, then on the role of market power.

Let me illustrate the role of negotiations with a comparatively uncontroversial example. Intel has a series of agreements with numerous OEMs about putting the Intel Inside and Centrino brands on their products to signal to users that the laptop includes a Wi-Fi compatible motherboard and antennae designed by Intel. In addition, Intel often includes certain compensation for the marketing expenses of putting the Intel copyrighted jingle inside a commercial for a PC.

This is a good example because this example illustrates that breakdowns can occur for many reasons. A few years ago Dell refused to carry the Centrino branded systems, and, accordingly, did not receive the quid pro quo compensation. Both parties went on their merry way for many years. Dell continued to carry both Intel products, but after that incident began to more prominently distribute designs with AMD chips. At the same time Intel reached deals with every other major OEM, and succeeded in making Centrino a feature of the majority of notebooks in use.

What does this example illustrate? First, that Intel's market power had its limits with Dell. It eventually reached a point in its negotiation with Dell where Intel gave Dell a take-it-or-leave-it offer and, indeed, Dell chose to leave it (unlike virtually everyone else in the industry). Second, as long as Dell had plenty of other options, the losses to Dell or society at large were not too large. Indeed, there might have been gains, since Dell's choices translated into more buyer options beyond the Centrino.

That last observation also relates to earlier remarks about the presence of alternatives. In the face of many alternatives, there are fewer policy justifications about negotiations between participants in an interdependent value chain. When there are plenty of suppliers, if users are unhappy with a supplier, or vendors are unhappy with a partner, they simply switch. In the presence of market power, however, these concerns are great.

Take-it-or-leave-it or refusal-to-deal can have some serious additional consequences when embedded in a network. Let me illustrate with one proposal. Some years ago there was a proposal to let all Internet participants simply negotiate compensation between them, so that Google/Yahoo/Disney would negotiate with Comcast/Time-Warner/Verizon, and every other possible combination. Intel's example suggests the problem with such a proposal: imagine the uproar among Internet users in the locations where such negotiations failed to come to resolution and no other close substitutes existed. It would be far worse than the brief uproar last year among Yankee fans who could not get local baseball telecasts due to a negotiation breakdown between Major Leagues Baseball, the Yankees, and a local cable provider. The most recent fight between Fox and cable firms is yet another example.

More to the point, it could be far worse for the Internet due to the way it potentially raises transaction cost for all other participants in the Internet value chain. If such an event were to be common place, imagine the uproar among Silicon Valley software developers and their VCs, who cannot be sure whether their software would deploy into a user base, or whether they would have to give up many of their innovative gains to those holding them up for a fee.

Imagine the uproar especially among providers of new innovative community software, peer-to-peer software, or other forms of communications software, who then would not sure one user can reach another in the near term. Most entrepreneurs have enough headaches to deal with and face low odds already; adding this additional one will only hurt commercialization of innovation.

As an aside, I have been surprised that its proponents have not also advocated mandatory binding arbitration. That is one mechanism that would reduce some of the transactions cost of such negotiations directly for the participants, and reduce some of the negative consequences for the transactions cost of other Internet participants. (It may be that such arbitration is not feasible in many cases because all relevant parties cannot be brought together, but this aspect deserves further attention.)

There is a more general point to make. In the context of a complex value chain and communications, the failure to make a deal in one part of the value has consequences for all others parts of the value chain and other users.

The presence of market power and lack of transparency adds to these concerns about negotiation. Firms with market power face incentives to not be transparent, or, as in a platform setting, at most, favor business partners over others in their actions related to transparency. Either case raises concerns about the raising of transaction cost to entrepreneurs and any participants in the Internet who gets treatment others than the favored partner.

This observation relates to one other aspect of the events surrounding Comcast, its apparent unwillingness or refusal to generate a transparent or predictable policy for accommodating some use of the application. Its policy effectively blocked use at arbitrary moments, and it offered no public explanation about why, though, after the fact, their motivation was rather obvious. All that was missing was a policy, for either senior management and/or technical staff, to make conduct clear to other participants in advance of that conduct.

By retaining all rights of discretion Comcast effectively made a take-it-or-leave –it-offer to both its own users and other providers of similar applications. It also retained the discretion to disable an innovative action by an entrepreneur.

To be clear, I am not focusing on Comcast in particular in order to comment on the FCC's decision in this case. Rather, I am using the example to illustrate the behavior that a firm with market

power could exhibit in its negotiation with other participants on the Internet and the consequences for the transaction cost of others. This example illustrates how lack of transparency, unilateral decision making, and no overt policy about such conduct could potentially raise transactions cost for others in the context of negotiations between providers.

More generally, due to the frequent change in technological opportunities, many of the factors in the Comcast event are not isolated or unique features of the event. In the near term, it is reasonable to expect negotiation to play a role in setting the conditions for the roll-out of further innovation. Many of the concerns about take-it-or-leave-it contracting or refusal-to-deal raise policy issues quite distinct from concerns (raised earlier) about opportunism and expropriation by dominant firms.

Bottom Line Summary

The commercial Internet differs from its predecessors. The setting today necessitates some active managerial role for carriers. Concerns about the use of managerial discretion arise in the presence of market power. Concerns also arise as a result of the interdependence between carriers and other participants in the Internet value chain. Wide participation in innovative activity does not necessarily arise in the absence of a minimal level of transparency and consistency by certain vendors, such as platform providers and vendors with market power.

These concepts and this model motivate a range of policy concerns about the transparency and consistency of firms with market power, such as carriers, and other platform providers. It also motivates policy efforts to improve transparency and consistency as the commercial Internet progresses and changes over the next decade.

Supporting material may be found in:

Shane Greenstein, 2009/forthcoming, "Glimmers and Signs of Innovative Health in the Commercial Internet," *Colorado High Technology Law Journal.*

Shane Greenstein, 2009/forthcoming, "Open Platform Development and the Commercial Internet," in (ed) Annabelle Gawer, *Platforms, Innovation, and Competition.* Pp 323-372.

Joel West & Siobhan O'Mahoney, 2008, "The Role of Participation Architecture in Growing Open Sourced Communities," 15, *Industry and Innovation*, April, at 145.

Copies of the first two papers can be found at:
http://www.kellogg.northwestern.edu/faculty/greenstein/images/research.html